NASA Technical Memorandum 88202

Ideal Shrinking and Expansion of Discrete Sequences

Andrew B. Watson

(NASA-TM-88202) IDEAL SHRINKING AND EXPANSION OF DISCRETE SEQUENCES (NASA) 36 p
CSCL 12A
N88-19100
Unclas
G3/59 0128720

January 1986

LIBRARY COPY
FEB 1
LANGLEY RESEARCH CENTER
LIBRARY, NASA
HAMPTON, VIRGINIA

National Aeronautics and
Space Administration

NASA Technical Memorandum 88202

Ideal Shrinking and Expansion of Discrete Sequences

Andrew B. Watson, Ames Research Center, Moffett Field, California

January 1986

National Aeronautics and
Space Administration

Ames Research Center
Moffett Field, California 94035

IDEAL SHRINKING AND EXPANSION OF DISCRETE SEQUENCES

Andrew B. Watson

Abstract-We describe ideal methods for shrinking or expanding a discrete sequence, image, or image sequence. The methods are ideal in the sense that they preserve the frequency spectrum of the input up to the Nyquist limit of the input or output, whichever is smaller. Fast implementations that make use of the discrete Fourier transform or the discrete Hartley transform are described. The techniques lead to a new multiresolution image pyramid.

This note describes a technique for shrinking and expansion of discrete sequences† which, though it follows rather directly from the nature of the discrete Fourier transform (DFT), appears not to be widely known. The two problems it solves are

(1) how to shrink a sequence of length N to length $M < N$ in such a way as to exactly preserve all the frequency content that can be preserved, and

(2) how to expand a sequence of length N to length $M > N$ in such a way as to exactly preserve the frequency content of the original and not to introduce frequencies not present in the original.

Sequence shrinking and expansion operations are widely used in digital signal processing and image processing, for example whenever the size or aspect ratio of an image must be changed. One recent development in the area of image analysis is the *image pyramid*, in which a single image is transformed into a number of copies that vary in resolution, each with a size proportional to resolution. [1] Although this transform expands somewhat the space required to store the image, the explicit representation of resolution information makes certain analyses of the image much easier. Image pyramids are generally constructed using shrinking and expansion operations, and we will show that the operations developed here lead to a new pyramid structure with some special virtues.

A large part of this paper consists of a detailed derivation of the ideal shrinking and expansion algorithms. So that the reader does not become lost in the details we

---

† We use the term *sequence* to refer to an ordered set of possibly complex numbers. As will be shown, the concepts generalize to two dimensions (images) and three dimensions (image sequences).

give here a brief synopsis. The basic idea is that a discrete sequence of length N contains frequencies up to $N/2$. These frequencies are represented in the discrete Fourier transform (DFT) of the sequence, which is also of length N. Within the DFT sequence there is a sub-sequence of length M which represents the frequencies up to $M/2$. If this subsequence is plucked from the DFT and inverse transformed, it will result in a new sequence of length M. This sequence will be an ideal low-pass filtered, subsampled version of the original. It will share all frequencies with the original up to $M/2$. We call it an "ideally shrunk" version of the original sequence. Similarly, an "ideally expanded" sequence is obtained by embedding the DFT of the original in a larger sequence of zeros, and inverse transforming. The only complexity (which is glossed over in the preceding synopsis) arises in the treatment of the Nyquist frequencies ($\pm N/2$ in the DFT of a sequence of length N), and we also provide an alternate, slightly-less-than-ideal algorithm which avoids these difficulties.

Since all shrinking and expansion operations involve filtering and sampling, we begin with a brief review of this subject in the domain of functions of a continuous variable.

1. Critical Sampling of Continuous Waveforms

In this discussion we follow Bracewell's use of the functions $\Pi(x)$, a unit pulse defined by

$$\begin{aligned} \Pi(x) &= 1 \qquad -1/2 < x < 1/2 \\ &= 1/2 \qquad x = -1/2 \text{ or } 1/2 \\ &= 0 \qquad \textit{elsewhere.} \end{aligned} \tag{1}$$

and $\Pi(x)$, an infinite train of unit impulses at unit intervals.

It is well known that a waveform band-limited below $w/2$ (the Nyquist frequency) can be exactly recovered from samples taken at intervals of $1/w$. [2, 3, 4] If it is necessary to sample an arbitrary waveform at frequency w, the portion of the spectrum between the Nyquist limits of $-w/2$ and $w/2$ can be preserved by removing the portion of the spectrum outside the Nyquist limits before sampling. This may be done by pre-filtering the waveform with an Ideal Low-Pass (ILP) filter whose Fourier transform is

$$L(u) = \Pi(u/w) \tag{2}$$

that is, a pulse of unit height and width w. The filtering is done by multiplying the Fourier transform of the waveform by $L(u)$ and inverse transforming the result. This is equivalent to convolving the waveform with $l(x)$, the inverse transform of $L(u)$, which is a sinc function. Following pre-filtering, sampling is accomplished by multiplying the waveform by $\Pi(xw)$, a train of impulses spaced at intervals of $1/w$. In the frequency domain, this replicates the prefiltered spectrum at intervals of the sampling frequency w. These replicas will overlap at $\pm w/2$. The even component at this frequency will be doubled in amplitude, while the odd component will cancel itself. So if the original spectrum is non-zero at $w/2$, then samples at a frequency of w preserve only the *even* portion of this component. This is known as *critical sampling*. [5]

2. Ideal Reconstruction of Continuous Waveforms

As noted, sampling replicates the band-limited spectrum at intervals of the sampling frequency w. To recover the band-limited spectrum it is necessary to multiply by a reconstruction filter $R(u)$,

$$R(u) = 1 \qquad -w/2 \leq u \leq w/2 \tag{3}$$
$$= 0 \qquad \textit{elsewhere.}$$

Note that this differs from $L(u)$ only at $\pm w/2$.

3. Discrete Fourier Transform

We assume the reader is familiar with the basic concepts of discrete complex sequences and of the discrete Fourier transform (DFT). We review several key points.

Let f_m be a complex sequence of length N. We define the Discrete Fourier Transform (DFT) of f_m as

$$F_k = \sum_{m=0}^{N-1} f_m \; W^{mk} \qquad k = 0, \cdots, N-1 \tag{4}$$
$$W = e^{-i\,2\pi/N}$$

The Inverse Discrete Fourier Transform (IDFT) is defined as

$$f_m = N^{-1} \sum_{k=0}^{N-1} F_k \; W^{-mk} \quad . \tag{5}$$

The relation between a sequence and its transform may be expressed by the notation

$$f_m \rightarrow F_k \tag{6}$$

which may be spoken as "f_m has DFT F_k". It should be noted that various definitions of the DFT exist in the literature, differing primarily in whether the factor N^{-1} is placed in the forward or the inverse transform. We follow that given in Rabiner, in Oppenheim and Schafer, and in Nussbaumer. [6, 7, 8]

4. Discrete Frequency Spectra

A sequence f_m of length N has a DFT F_k of length N, which defines its discrete frequency spectrum. By convention, the indices $k = 0, \cdots, N-1$ of the transform

map to frequencies $w = -N/2, \cdots, N/2-1$ according to the rule

$$w = ((k - N/2) \textit{ modulus } N) - N/2 \tag{7}$$

as shown in Fig. 1 Where the DFT is concerned, a sequence and its transform are regarded as periodic: for a sequence f_m of length N, $f_{m-kN} = f_m$ for integer k. Note that for a sequence of length N, the "highest" frequency present is $-N/2$, (the Nyquist frequency). Since the spectrum is periodic with period N, the component at $N/2$ is the same as the component at $-N/2$. Hence a sequence of length N contains only an even component at the Nyquist frequency.

8. Ideal Re-Sampling

To retain a close analogy to the continuous case, we have interpreted *sampling* of a discrete sequence to mean multiplying by a sampling sequence (zeroing certain elements), but not changing its length. We will use the term *re-sample* to describe an operation on a discrete sequence which does change its length.

When a discrete sequence is re-sampled some of its properties will change and some will remain the same. Choice of a re-sampling scheme must be guided by which properties the user wishes to preserve. In the scheme discussed here, we preserve, where

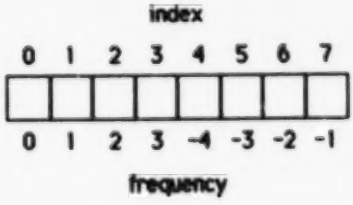

Figure 1. Arrangement of frequencies in the complex DFT.

possible, the portions of the frequency spectrum shared by old and new sequences. We call this *ideal re-sampling* (IR), so-called by analogy to an ideal low-pass filter.

The following discussion is illustrated in Fig. 2. We begin with a discrete sequence f_m of length N with DFT F_k (Fig.s 2a and b). We wish to resample f_m to a sequence of length $M < N$. It should be clear that if N/M is not an integer, then the new sample points will "fall in between" the samples in the original sequence. Thus it becomes necessary to convert the discrete sequence to a function of a continuous variable before sampling, and then back to a discrete sequence after sampling.

To do this, we construct a function

$$f'_N(x) = \sum_{-N/2}^{N/2-1} f_m \, \delta(x - m/N) \tag{8}$$

that is a set of N impulses spaced at intervals of $1/N$, multiplied by the values of the discrete sequence (Fig. 2c).

We can do likewise with the DFT F_k (Fig. 2d), except that the impulses are spaced at unit intervals, and the result is expressed as a function of a frequency variable u,

$$F'_N(u) = \sum_{-N/2}^{N/2-1} F_k \, \delta(u - k) \quad . \tag{9}$$

Next we convolve f'_N with $\Pi(x)$,

$$f_N(x) = f'_N(x) * \Pi(x) \quad . \tag{10}$$

This has the effect of replicating the set of impulses at unit intervals. Since the set extends over a distance of $(N-1)/N$, this results in a periodic function (Fig. 2e).

Likewise, F'_N can be convolved with $\Pi(u/N)$ (a pulse train with intervals of N) to replicate the set of impulses at intervals of N,

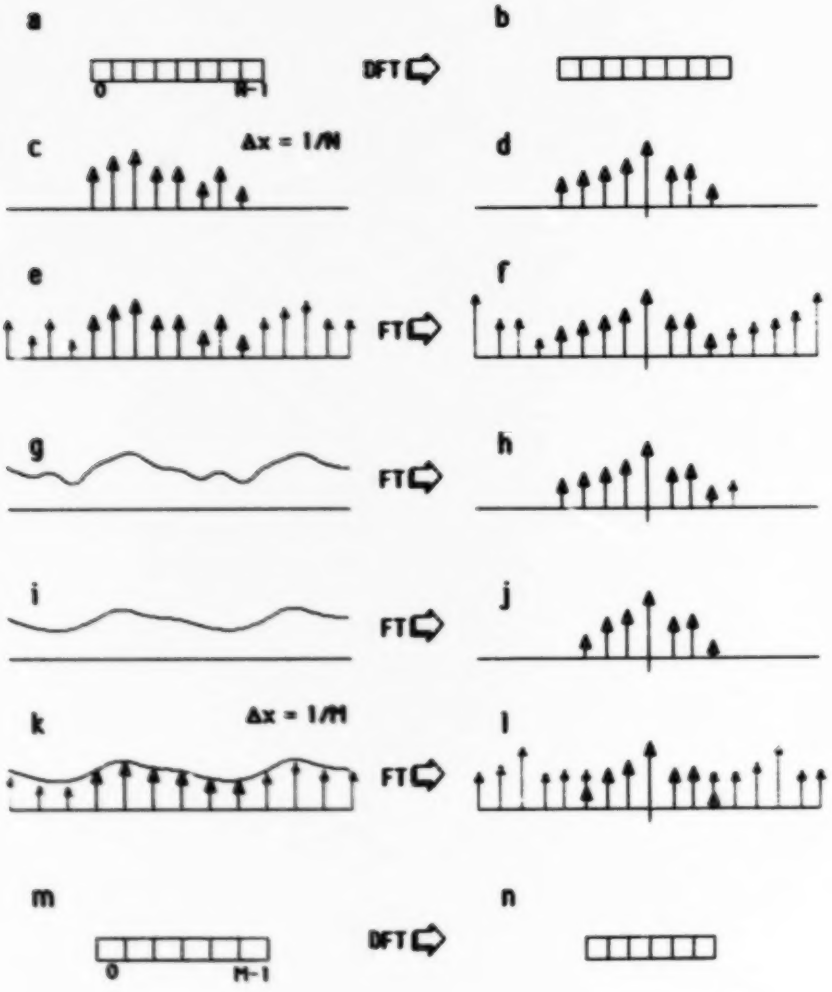

Figure 2. Ideal shrinking of a discrete sequence of length N to a sequence of length $M < N$.

$$F_N(u) = F'_N(u) * \Pi(u/N) \quad . \tag{11}$$

as shown in Fig. 2f.

Now we will show, as indicated in Fig. 2, that $F_N(u)$ is the Fourier transform of $f_N(x)$. Making use of the convolution theorem and the Fourier transform of an impulse, and the fact that Π is its own Fourier transform,

$$FT\{f_N(x)\} = \left[\sum_{-N/2}^{N/2-1} f_m \; e^{-i\,2\pi um/N} \right] \Pi(u) \quad . \tag{12}$$

Note that, whatever the function within the brackets, it will be multiplied by $\Pi(u)$ and so only its values at integer values of u will matter. But when the integer variable k is substituted for u into the bracketed expression, it becomes an exact expression for the DFT of f_m, that is, F_k,

$$FT\{f_N(x)\} = \left[F_k \right]_{k=u} \Pi(u) \quad . \tag{13}$$

But, due to the cyclic nature of F_k, this is an expression for $F_N(u)$, thus

$$FT\{f_N(x)\} = F_N(u) \tag{14}$$

The above steps provide a general scheme for moving between a sequence and its DFT, and a function and its FT.

The function $f_N(x)$ corresponds uniquely to a continuous, band-limited function $f(x)$ which may be thought of as the function from which $f_N(x)$ was sampled. As noted above, $f(x)$ can be exactly reconstructed by multiplying the FT of $f_N(x)$, that is, $F_N(u)$, by the function $R(u/N)$. As shown in Fig. 2i, this has the effect of stripping away the replicas of the original set of impulses, and returning to the function of Fig. 2d, plus an extra impulse at $N/2$. The corresponding operation in the space domain is convolution with the function $\mathrm{sinc}(Nx)$, which converts the set of impulses

into the continuous periodic function shown in Fig. 2h.

In the next step, we wish to sample the continuous function at a rate of $1/M$ so that each period will be described by exactly M samples. As noted above, we can do this without aliasing if the function is first ILP filtered. This can be done by multiplying the transform $F(u)$ by the function $\Pi(u/N)$,

$$G(u) = \Pi(u/M)\, F(u) \tag{15}$$

As shown in Fig. 2j, this reduces the number of impulses to $M+1$, and in addition halves the values of the two outermost impulses at $M/2$ and $-M/2$. In the space domain, the resulting function $g(x)$ is somewhat smoother than $f(x)$. It could have been obtained from $f(x)$ by convolution with the function $M\ \mathrm{sinc}(Mx)$.

Now we can sample $g(x)$ at intervals of $1/M$ by multiplying by $\Pi\Pi(Mx)$ (Fig 2k),

$$g_M(x) = g(x)\, \Pi\Pi(Mx) \tag{16}$$

But this is equivalent, in the frequency domain, to convolution with the function $\Pi\Pi(u/M)$ (Fig. 2l),

$$G_M(u) = G(u) * \Pi\Pi(u/M) \tag{17}$$

As illustrated, this convolution replicates the transform $G(u)$ at intervals of M. This causes the left and rightmost impulses in the original transform to overlap their opposite numbers in the replicas. Thus $G_M(u)$ is periodic with period M, and the value at $-M/2$ is equal to the value at $M/2$, which in turn is equal to the average of the values at these two frequencies in the original DFT F_k.

The final step is to reverse the process whereby we converted a discrete sequence and its DFT into a function and its FT. The discrete sequence g_m is taken as the amplitudes of the impulses making up one period of the function $g_M(x)$, and likewise

G_k is drawn from one period of $G_M(u)$. Thus we arrive at the sequence g_m and its DFT G_k, both of length M. Note that the components of the sequence G_k are identical to the corresponding components in F_k, save for that at the new Nyquist frequency $M/2$, where only the even component of the original sequence is preserved. Thus the sequence g_m is an ideal re-sampling of f_m.

It should be clear from the proceeding that the ideal resampling operation is much simpler in the frequency domain than in the space domain. We make this explicit below in the form of ideal shrinking and expansion algorithms.

6. Ideal Shrinking

Let M be an integer less than N and let Q be the rational number M/N. We define $ISHRINK_Q$ (Ideal Shrink) as the operation that shrinks a sequence f_m of length N to a sequence g_m of length $M = QN$, by the rule

$$g_m = ISHRINK_Q\{ f_m \} = IDFT\{TRIM_Q\{ DFT\{f_m\}\}\} \tag{18}$$

The $TRIM_Q$ operator acts on $F_k = DFT\{f_m\}$ to produce $G_k = DFT\{g_m\}$ according to the rule

$$G_k = F_k \qquad k = -M/2+1, \cdots, M/2-1 \tag{19}$$

$$G_{M/2} = (F_{M/2} + F_{-M/2})/2$$

Frequencies below the new Nyquist are drawn directly from corresponding frequencies in the old sequence. The single value *at* the new Nyquist is taken as the average of the two corresponding frequencies in the old spectrum. One way of envisioning the action of the $TRIM_Q$ operator is illustrated in Fig. 3. The two ends of the transform, up to each new Nyquist, are pulled off. The Nyquists are halved, and the two pieces are put

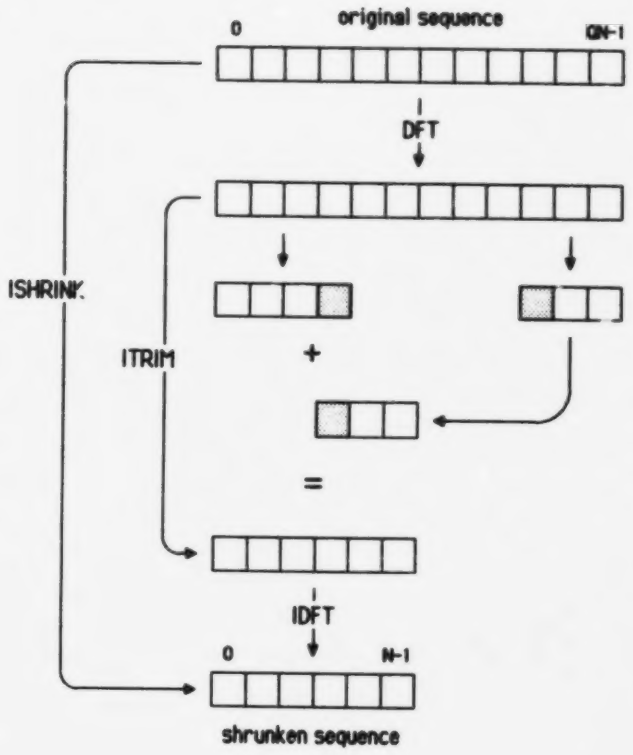

Figure 3. The $ISHRINK_Q$ and $TRIM_Q$ operators. Shaded elements are multiplied by 1/2.

together so that the Nyquists add.

7. Ideal Expansion

The logic of ideal expansion is nearly identical to that of ideal shrinking. As shown in Fig. 4, it differs only at two steps. First, when the continuous waveform $f(x)$ is ILP filtered by multiplying its transform $F(u)$ by $\Pi(u/M)$, this has no effect, because the spectrum has no components at or beyond the new Nyquist limits ($u \geq M/2$). Thus the waveforms and spectra in Fig.s 2g-j are identical. Second, when the function $g(x)$ is sampled, its spectrum is replicated at intervals of M, leaving a gap between the end of the original spectrum and the start of the first replica. Note that the values of the original spectrum are preserved in the resampled spectrum, thus the operations pictured in Fig. 3 are an ideal resampling of the sequence f_m. These observations can be condensed into an ideal expansion algorithm.

Let f_m be a sequence of length N, with DFT F_k. We define $IEXPAND_Q$ as the operator that expands f_m to a sequence g_m of length M according to the rule

$$IEXPAND_Q \{ f_m \} \rightarrow PAD_Q \{ F_k \} \tag{20}$$

As in the case of ideal shrinking, the expansion is achieved by way of an operation on the transform, which we give the name PAD_Q, and which is defined by

$$\begin{aligned} G_k &= F_k \qquad k = -N/2, \cdots, N/2 \\ &= 0 \qquad \textit{elsewhere}. \end{aligned} \tag{21}$$

Frequencies below the old Nyquist are mapped directly into corresponding frequencies in the new sequence, and the remainder of the spectrum is filled with zeros.

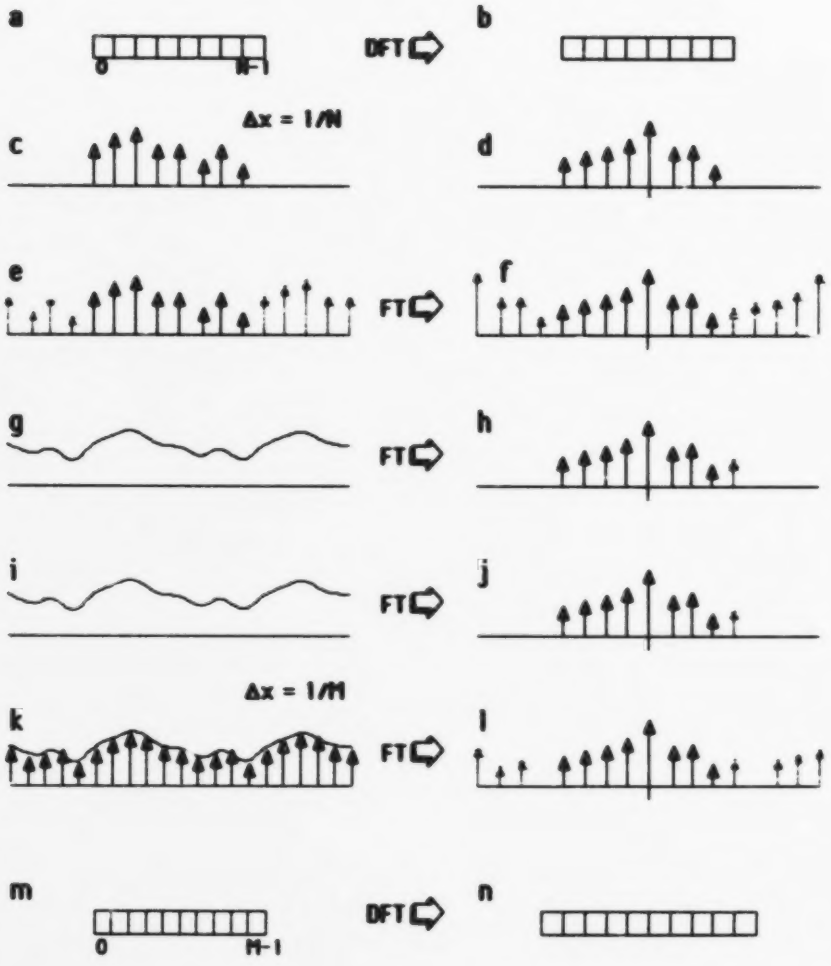

Figure 4. Ideal expansion of a discrete sequence of length N to a sequence of length $M > N$.

It can be shown that if the sequences shown in Fig. 2 are now regarded as purely real DHT's (rather than complex DFT's), then all of the steps of the derivation of the ideal shrink operation are equally valid. Thus somewhat simpler versions of the $TRIM_Q$ and PAD_Q can be constructed in which real DHT's are padded or trimmed, and inverse DHT'd, to obtain ideally shrunk or expanded sequences.

11. The Ideal Pyramid

In the area of image processing and analysis there has been some interest in "pyramid" schemes. [12, 13, 14, 15, 1] Such schemes typically subject an image to repeated low-pass filtering and sub-sampling, each cycle creating one layer of the pyramid. This low-pass pyramid can then be used to create a band-pass pyramid, by differencing adjacent low-pass layers. [13, 14] The shrink and expand operators introduced here lead directly to a pyramid structure. Although pyramid algorithms are typically used on two-dimensional images, for simplicity we begin with a one dimensional example. Let x_0 be an input sequence of length NS^I, where N, S, and I are integers. The original image is considered the base of the pyramid. Let $Q = 1/S$. Each of the other I layers of the low-pass pyramid is given by

$$x_i = ISHRINK_Q \{ x_{i-1} \} \quad i = 1, \cdots, I \tag{20}$$

In words, each successive layer is obtained by shrinking the previous layer by a factor $Q = 1/S$. Thus layer i is of length NS^{I-i}. In particular, the first layer is of size NS^I, and the last layer is of size N.

Each layer of the band-pass pyramid is created by subtracting from each lowpass layer an expanded version of the next layer,

PRECEDING PAGE BLANK NOT FILMED

15, 16, [illegible], 8

$$y_i = x_i - IEXPAND_Q \{ x_{i+1} \} \quad i = 0, \cdots, I-1 \tag{30}$$

The last band-pass layer y_I is taken as equal to the last lowpass layer x_I. When $Q = 1/2$, each band-pass layer has a one octave bandwidth. Examples of low-pass and band-pass ideal pyramids are shown in Fig. 7. Also shown are versions of the low- and band-pass pyramids expanded up to full size by means of the $IEXPAND_Q$ operator. Note that the expanded band-pass waveforms add up to form the original sequence. Creation of the pyramid requires only one forward DFT. From that point on only element selection and comparatively small IDFT's are required.

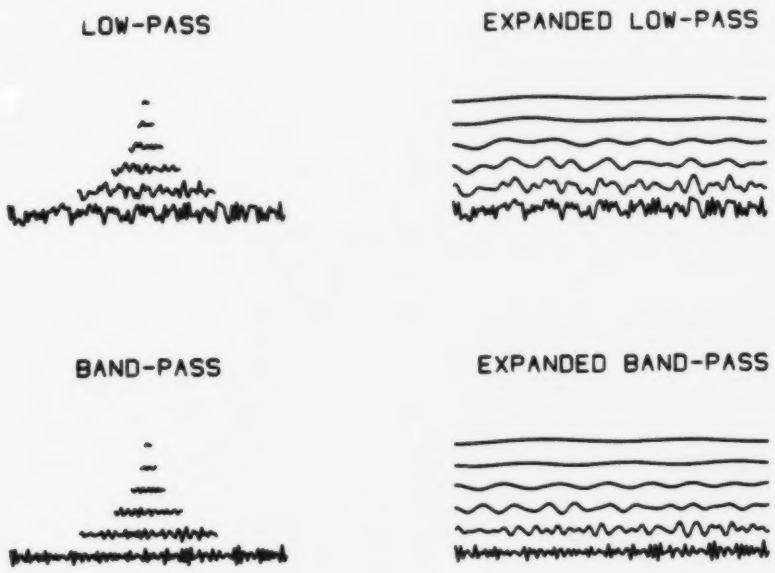

Figure 7. Ideal 1D pyramids. The original is a random sequence of length 128. $Q = 1/2$.

12. Ideal Expansion and Shrinking in Two Dimensions

The two-dimensional version of the shrink algorithm is routine. Since the two-dimensional DFT is a one-dimensional DFT of each row (or column) followed by a one-dimensional DFT of each column (or row), the Trim algorithm can be applied in sequence to each row, then each column (or vice versa) as diagrammed in Fig. 8. A more direct approach is diagrammed in Fig. 9. The horizontal and vertical Trim operations have been combined. It remains only to shift and add the selected regions of the transform. The lower panel of Fig. 9 shows the Trim operation following a remapping of the transform indices to place the origin near the center. This makes it clearer that the Trim operator is selecting the low-pass core of the transform. Figure 10 is a flow diagram of the steps involved in the creation of an ideal band-pass pyramid, and of the reconstruction of the original image from the pyramid. Figures 11-13 show the use of the 2D $ISHRINK_Q$ and $IEXPAND_Q$ operations to create an image pyramid.

13. Ideal Expansion and Shrinking in Three Dimensions

Generalization of the $ISHRINK_Q$ and $IEXPAND_Q$ operators to three or more dimensions is straightforward, since they can always be expressed as a sequence of one-dimensional Trim operations. When the third dimension is time, the input is an image sequence, and the $ISHRINK_Q$ and $IEXPAND_Q$ operators enlarge or reduce the duration of the sequence.

13.1. Slow Motion

As an example of the application of the $IEXPAND_Q$ operation in the time domain, consider the problem of creating a slow-motion sequence. Suppose an image sequence is

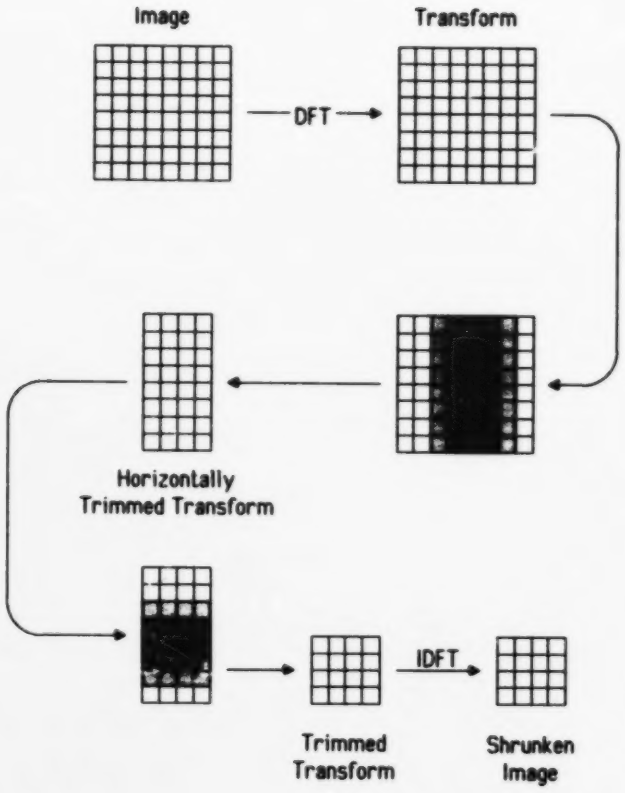

Figure 8. The Shrink algorithm in two dimensions. The one-dimensional Trim algorithm is applied first to the rows, then to the columns.

ORIGINAL PAGE IS
OF POOR QUALITY

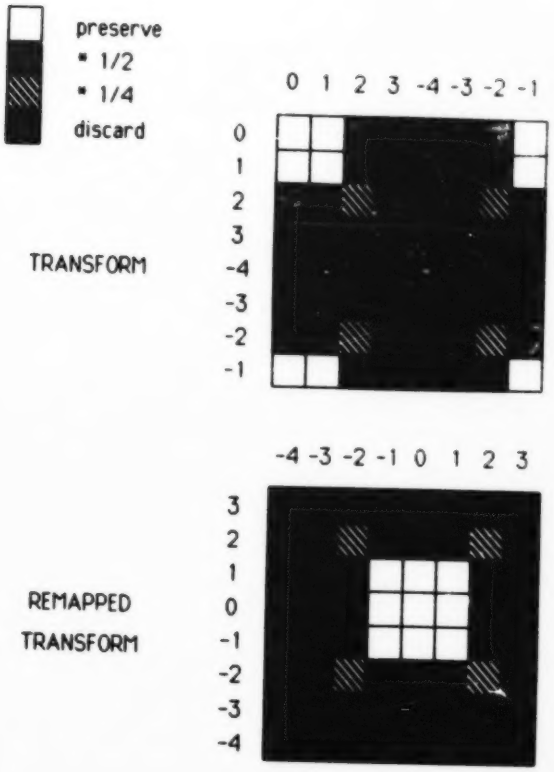

Figure 9. The two-dimensional Trim operation. The upper panel shows the conventional ordering of elements in the DFT. The lower panel shows the result of shifting the transform by 4 in both dimensions so as to place the origin near the center. The algorithm would be completed by extracting the leftmost column of the non-zero area and adding it to the leftmost column, then extracting the uppermost row and adding it to the lower-most row.

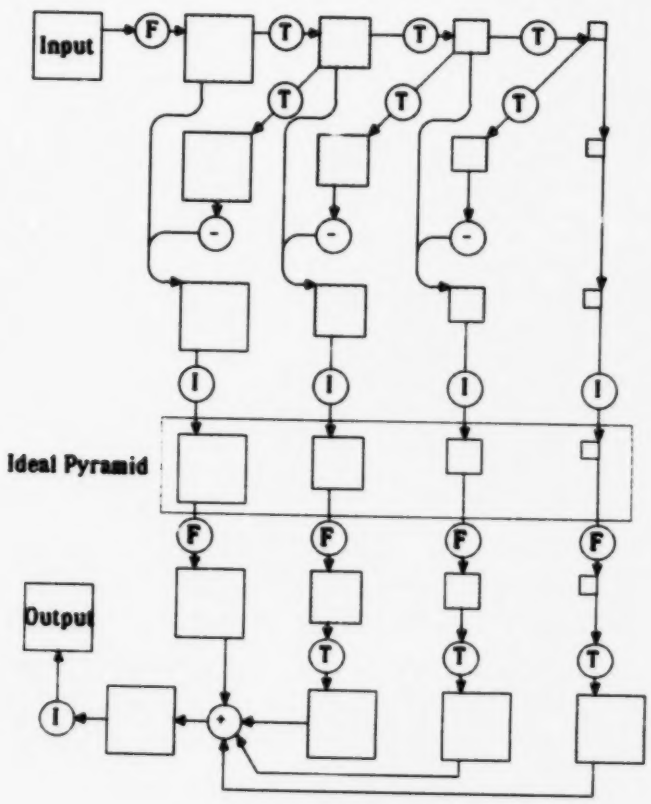

Figure 10. Flow diagram for the creation and reconstruction of an ideal band-pass pyramid. The letter meanings are: F, forward DFT; I, inverse DFT; T, Trim operation to enlarge or shrink a transform. The stippled images are DFTs.

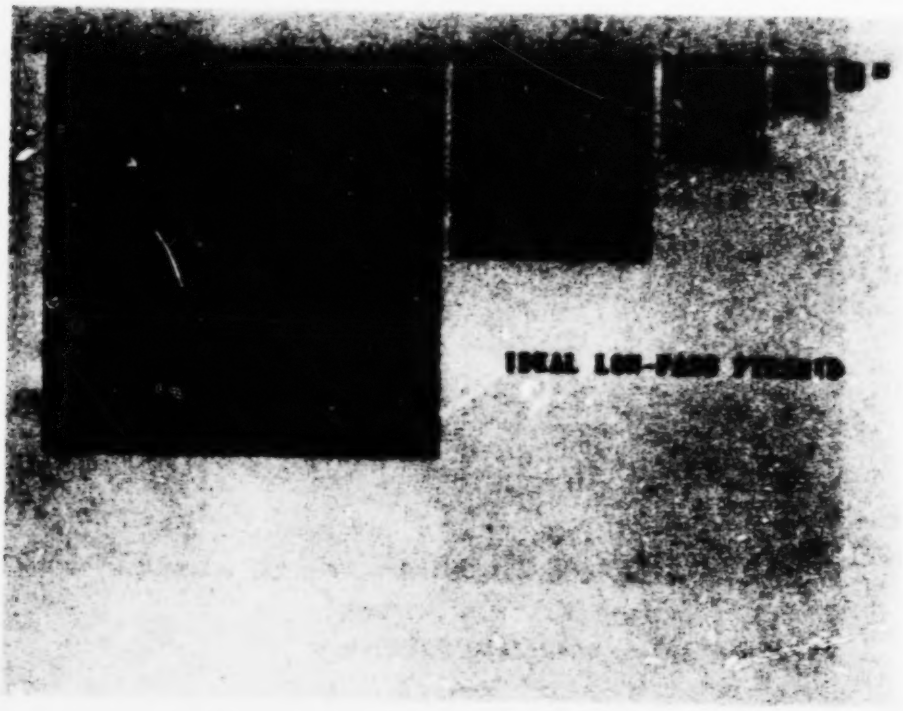

Figure 11. An Ideal Low-Pass Image Pyramid created by ideal shrinking and expansion. The value of Q is 1/2.

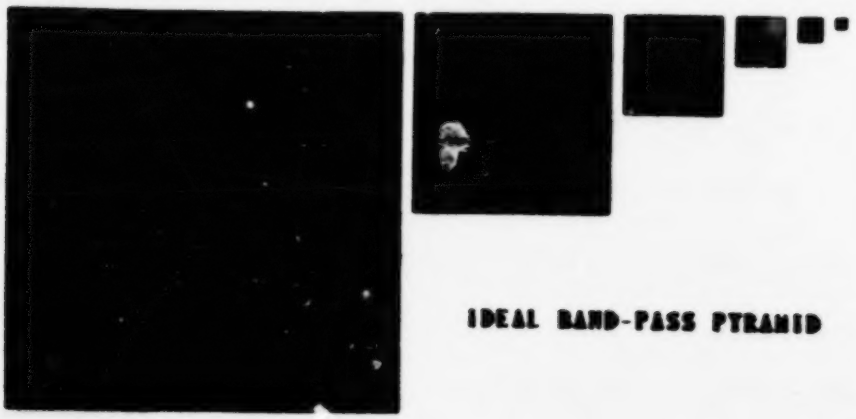

Figure 12. An Ideal Band-Pass Pyramid. The value of Q is 1/2, so each band-pass image has a one octave bandwidth.

ORIGINAL PAGE IS
OF POOR QUALITY

Figure 13. Expanded Ideal Low-Pass and Band-Pass pyramids. Each picture in the second row is the difference between the picture directly above and the picture above and to the right. The sum of the pictures in the bottom row equals the image in the upper left.

recorded at 60 Hz for 1 second. Normal playback would be at a rate of 60 Hz. If played back at 6 Hz (one frame every 167 msec) the sequence would last 10 seconds and motion of the imagery would be slowed by a factor of 10. But there would be large jumps between frames (which we call "jerkies", by analogy to the "jaggies" in spatial images). One solution is to record the original at 600 Hz, and play back at 60 Hz, but this is usually quite impractical and invariably expensive. A better solution is to use the $IEXPAND_Q$ routine in the time domain, with Q=10. The resulting imagery would be slowed, but would have no jerkies.

13.3. In-Betweening

Animation is typically done by creating a sequence of "key frames", and the creating the intervening sequences by "in-betweening". Automatic methods of in-betweening

ORIGINAL PAGE IS
OF POOR QUALITY

are being sought. The $IEXPAND_Q$ routine may provide such a method. A set of key frames are regarded as a sequence which is expanded to the desired length. Each key frame will fade as the next key frame grows in contrast. While objects will not in general be drawn correctly in their intermediate positions, they may nonetheless appear subjectively to occupy these intermediate positions.

13.3. 3D Image Compression

Burt and Adelson have shown the usefulness of pyramid schemes as a tool for image compression. It seems likely that the ideal pyramid would provide similar benefits. Inclusion of the time dimension in the pyramid makes it possible to remove additional redundancies in the time domain, as well as to take advantage of the spatio-temporal tuning of the human visual system, to reduce still further the bit rate for dynamic imagery.

14. A simpler, sub-ideal algorithm

The only complexity in the $ISHRINK_Q$ algorithm arises in the treatment of the Nyquist frequencies, which must be averaged to yield their even component. While this step is simple in one dimension, it becomes somewhat complicated in two and three dimensions. Care must be taken that the resulting Nyquist frequencies are indeed the result of the sequence of separate one-dimensional $ISHRINK_Q$ operations. An alternative is to discard the Nyquist frequencies, in which case the Trim algorithm consists of selecting only those frequencies *below* the new Nyquist, and inserting a zero at the Nyquist. Often the even Nyquist component in an image can be removed without substantially degrading its visual appearance. Furthermore, when constructing a pyramid,

those components omitted from one layer are captured by another, so this simplified algorithm just pushes the even Nyquist component in each layer down to the next layer.

15. Comparison with other pyramid schemes

It is interesting to compare the three shrinking methods that have been mentioned. Each has its vices and virtues, a brief tally of which is given in Table 1.

An important practical difference is that the $ISHRINK_Q$ method allows the output to be any integer size, while the other two methods require the output size to be smaller than the input by some integer factor. This is because the latter two methods resample the image on the same lattice of sample points as are used in the original (skipping every 1/Q-1 samples), while the $ISHRINK_Q$ method in effect creates new sample points.

Table 1. Vices and virtues of three image-shrinking methods.

Feature	Quality	Method		
		Block	Laplacian	Ideal
radial symmetry	good	no	yes	no
separable	good	yes	yes	yes
small kernel	good	yes	yes	no
aliasing	bad	yes	yes	no
attenuation	bad	yes	yes	no
overlap in frequency	bad	yes	yes	no
non-integer size ratio	good	no	no	yes
total score		-3	-1	3

All three shrinking schemes consist of some filtering followed by re-sampling, but differ in the filters used. In the method of Tanimoto and Pavlidis [12] the sequence is convolved with a rectangular pulse, or equivalently, its transform is multiplied by a sinc function. In the Ideal Pyramid, the sequence is convolved with a sinc function, or equivalently, its transform is multiplied by a rectangular pulse. The filters used in the two schemes are thus DFT's of one another, and are at opposite poles of a spectrum of algorithms. The Gaussian filter used by the DOLP and Laplacian pyramids [15, 13] may be regarded as intermediate between these two.

The Gaussian is the only 2D function that is both radially symmetric and separable. The pulse and Gaussian filters have the virtue that they have small kernels in the space domain, which means that it is practical to compute them by direct convolution using integer arithmetic. They both suffer in failing to precisely bandlimit the signal before re-sampling, and thus both exhibit aliasing. Both also attenuate signals at frequencies below the Nyquist, which is generally manifest as more blurring than is strictly necessary in the shrunken sequence. A possibly attractive property found only in the sinc filter is that the resulting band-pass sequences have spectra that do not overlap (more precisely, overlap only in the even component at the Nyquist frequency). They thus provide a cleaner partition of the image into resolution-specific sub-images, should that be found useful. This list is not exhaustive, and the values associated with each of these properties will depend strongly upon the application. In the table, the total value assigned to each method is clearly not a statement regarding the ultimate value of each technique, and is given only to save the reader the effort of this calculation.

Acknowledgements:

I thank Al Ahumada and Denis Pelli for many useful comments, and Brian Wandell for drawing my attention to the Hartley transform.

References

[1] A. Rosenfeld *Multiresolution Image Processing and Analysis.* New York: Springler-Verlag, 1984.

[2] H. Nyquist, "Certain topics in telegraph transmission theory," *Transactions A.I.E.E.*, , pp. 617-644, 1928.

[3] C. E. Shannon, "Communication in the presence of noise," *Proceedings IRE*, vol. 37, pp. 10-21, 1949.

[4] A. J. Jerri, "The Shannon sampling theorem - Its various extensions and applications: A tutorial review," *Proc. IEEE*, vol. 65, pp. 1565-1596, 1977.

[5] R. N. Bracewell *The Fourier transform and its applications.* New York: McGraw-Hill, 1978.

[6] IEEE Acoustics, Speech, and Signal Processing Committee *Programs for digital signal processing.* New York: IEEE Press, 1979.

[7] H. J. Nussbaumer *Fast Fourier transformation and convolution algorithms.* New York: Springer-Verlag, 1982.

[8] A. V. Oppenheim and R. W. Schafer *Digital signal processing.* Englewood Cliffs, N.J.: Prentice Hall, 1975.

[9] R. N. Bracewell, "Discrete Hartley transform," *Journal of the Optical Society of America*, vol. 73, pp. 1832-1835, 1983.

[10] R. N. Bracewell, "The fast Hartley transform," *Proceedings of the IEEE*, vol. 72, pp. 1010-1018, 1984.

[11] A. B. Watson and Allen Poirson, "A separable two-dimensional discrete Hartley transform," *Journal of the Optical Society of America A*, vol. submitted, 1985.

[12] S. Tanimoto and T. Pavlidis, "A Hierarchical data structure for picture processing," *Computer Graphics and Image Processing*, vol. 4, pp. 104-119, 1975.

[13] J. L. Crowley and R. M. Stern, "Fast computation of the difference of low-pass transform," *IEEE Transactions on pattern analysis and machine intelligence*, vol. 6, pp. 212-222, March 1984.

[14] P.J. Burt, "Fast filter transforms for image processing," *Computer Graphics and Image Processing*, vol. 16, pp. 20-51, 1981.

[15] Peter J. Burt and Edward H. Adelson, "The laplacian pyramid as a compact image code," *IEEE Transactions on Communications*, vol. COM-31, pp. 532-540, 1983.

1. Report No.	2. Government Accession No.	3. Recipient's Catalog No.
NASA TM-88202		

4. Title and Subtitle	5. Report Date
IDEAL SHRINKING AND EXPANSION OF DISCRETE SEQUENCES	December 1985
	6. Performing Organization Code
7. Author(s)	8. Performing Organization Report No.
Andrew B. Watson	A-86070
	10. Work Unit No.
9. Performing Organization Name and Address	
Ames Research Center Moffett Field, CA 94035	11. Contract or Grant No.
	13. Type of Report and Period Covered
12. Sponsoring Agency Name and Address	Technical Memorandum
National Aeronautics and Space Administration Washington, DC 20546	14. Sponsoring Agency Code 505-47

15. Supplementary Notes

Point of Contact: Andrew B. Watson, Ames Research Center, MS 239-3, Moffett Field, CA 94035 (415)694-5419 or FTS 464-5419

16. Abstract

We describe ideal methods for shrinking or expanding a discrete sequence, image, or image sequence. The methods are ideal in the sense that they preserve the frequency spectrum of the input up to the Nyquist limit of the input or output, whichever is smaller. Fast implementations that make use of the discrete Fourier transform or the discrete Hartley transform are described. The techniques lead to a new multiresolution image pyramid.

17. Key Words (Suggested by Author(s))	18. Distribution Statement
Image processing Sampling Fourier transform Image pyramid Decimation	Unlimited Subject category - 59

19. Security Classif. (of this report)	20. Security Classif. (of this page)	21. No. of Pages	22. Price*
Unclassified	Unclassified	36	A03

*For sale by the National Technical Information Service, Springfield, Virginia 22161

685-297/58131